

Optimizing positional scoring rules for rank aggregation*

Ioannis Caragiannis
George A. Krimpas

Xenophon Chatzigeorgiou
Alexandros A. Voudouris

Department of Computer Engineering & Informatics
University of Patras, 26504 Rion, Greece

Abstract

Nowadays, several crowdsourcing projects exploit social choice methods for computing an aggregate ranking of alternatives given individual rankings provided by workers. Motivated by such systems, we consider a setting where each worker is asked to rank a fixed (small) number of alternatives and, then, a positional scoring rule is used to compute the aggregate ranking. Among the apparently infinite such rules, what is the best one to use? To answer this question, we assume that we have partial access to an underlying true ranking. Then, the important optimization problem to be solved is to compute the positional scoring rule whose outcome, when applied to the profile of individual rankings, is as close as possible to the part of the underlying true ranking we know. We study this fundamental problem from a theoretical viewpoint and present positive and negative complexity results and, furthermore, complement our theoretical findings with experiments on real-world and synthetic data.

1 Introduction

Social choice theory [4] studies voting rules (also known as social choice or social welfare functions) that compute a winning alternative or a ranking of the available alternatives from voter preferences. Typically, the preference of each voter is supposed to be a ranking over *all* available alternatives. We deviate from this assumption and, instead, we focus our attention to settings in which each voter (or, better, agent for our purposes) ranks only a small subset of the alternatives. Such *incomplete* rankings seem to be non-standard in the literature; [7, 8, 15] are some notable exceptions.

The setting we have in mind is motivated by crowdsourcing [10] and rating applications. For example, assume that a requester would like to rank a huge set of alternatives using expert opinions from a crowd of workers. Asking each worker for her opinion on the whole set of alternatives (i.e., for a full ranking) would rather result in poor information. Most probably, the worker will not be aware of most of the alternatives. Even if she tries to obtain additional information, coming up with consistent comparisons between alternatives that she knows well and alternatives that she has no idea about, would be rather impossible, given their huge number. Instead, this task would be much easier if workers focused on small sets of alternatives. The requester could give each worker a different set of few alternatives to rank. Then, processing smaller inputs would be easier for the requester as well.

This approach has been recently exploited in the context of ordinal peer grading in MOOCs; see [5, 6, 14, 16] for approaches of this flavour. In such settings, the task of grading an exam with many participating students is outsourced to the students themselves. Each student is given a small number

*This work was partially supported by the Caratheodory research grant E.114 from the University of Patras and by a Ph.D. Scholarship from the Onassis Foundation. Emails: {caragian, chatzigeorgiou, krimpas, voudouris}@ceid.upatras.gr.

of exam papers to rank and the final grading (a ranking of all students) is obtained by aggregating the inputs provided by the students.

In a rating application we envision, users of a hotel booking system are asked to rank hotels they have stayed recently in a specific city and the rating application aims to compute a full ranking of the hotels (or, possibly, different rankings for different relevant criteria). Clearly, each user can provide meaningful feedback for just a few hotels. Again, in this scenario, the system might ask each user to focus only on a subset of the hotels she knows.

Besides the different sets of alternatives each individual is asked to rank in the above scenarios, another implicit feature is that there is an *underlying true ranking* of all alternatives (e.g., the ranking of exam papers in terms of their quality or the ranking of hotels in terms of their facilities) that we would like to compute when aggregating the individual preferences. Can we do so using simple voting-like rules? We follow an optimization approach which can be described with the following question: Assuming that we have partial knowledge of the underlying true ranking and access to sampled profiles, which is the rule that yields an outcome that is as consistent as possible to (our partial knowledge of) the underlying true ranking when applied to the sampled profiles?

We study the above question for *positional scoring rules* (or, simply, scoring rules), which have played a central role in social choice theory. Two factors that have led to this decision are their simplicity and effectiveness; simplicity follows by their definition and effectiveness is justified by our experimental results. In particular, we consider settings in which each agent is asked to rank the same number d of alternatives. A positional scoring rule in our setting is defined by a scoring vector (s_1, s_2, \dots, s_d) . It takes as input the incomplete individual rankings of the agents and computes scores for alternatives as follows. An alternative gets s_k points each time it is ranked k -th by an agent and its score is its total number of points. The final ranking is obtained by ordering all alternatives in terms of their scores, in non-increasing order. Now, a profile of individual incomplete rankings and desired relations for pairs of alternatives (to be thought of as parts of the underlying true ranking) with corresponding weights indicating the importance of each relation, we would like to compute the positional scoring rule whose outcome when applied on the profile maximizes the total weight of the desired pairwise relations it satisfies. This is related to learning-theoretic studies where a scoring rule that as consistent as possible to given examples is sought; e.g., see the paper by Boutilier et al. [2] and Procaccia et al. [13]. The main difference of the current paper (besides our assumption on profiles with incomplete rankings) is its optimization flavour. We refer to this seemingly fundamental optimization problem as **OptPSR**.

Our technical contribution is as follows. We present an exact algorithm that solves **OptPSR** in time that depends exponentially only on the parameter d (Section 3). Hence, our algorithm runs in polynomial time when d is constant. For instances with high values of d , we show that a simple t -approval voting rule (that uses the scoring vector with t 1s followed by $d - t$ 0s) yields a $1/d$ -approximate solution. We show that this bound is tight by constructing an instance in which any approval voting rule is $1/d$ -approximate. We prove that **OptPSR** is hard to approximate and present an explicit inapproximability bound of $23/24$. This result follows by an approximation-preserving reduction from the problem MAX-3LIN-2 of maximizing the number of satisfied equations in an over-determined system of linear equations modulo 2 and exploits a famous inapproximability result due to Håstad [9]. These results can be found in Section 4. In Section 5, we describe experiments on real-world and synthetic profiles. Our experimental results show that scoring rules perform remarkably well and recover almost 100% of the desired constraints in many interesting scenarios; this justifies our choice to study scoring rules (and the optimization problem **OptPSR**) in the first place. We begin with preliminary definitions in Section 2 and conclude with open problems in Section 6.

2 Problem statement

We consider settings with a set of *agents* N and a set of *alternatives* A . Agent i expresses her preference over a subset $A_i \subseteq A$ of alternatives; her preference is a ranking of the alternatives in A_i . A preference profile (or simply, a *profile*) consists of the preferences of all agents. In this work, we assume that all agents have the *same* number $d \geq 2$ of alternatives in their preference, i.e., $|A_i| = d$ for each agent i .

A social welfare function takes as input a profile Π and it outputs a ranking of all alternatives in A . A positional scoring rule (or, simply, a scoring rule) is a social welfare function that uses a scoring vector $\mathbf{s} = (s_1, \dots, s_d)$ with $s_i \geq s_{i+1}$ for $i = 1, \dots, d-1$ and $s_d \geq 0$; the alternative at position k in each vote is assigned s_k points and the ranking of the alternatives is produced by ordering them in monotone non-increasing order in terms of their total points (or score). Formally, for an alternative x , let $\nu_j(x, \Pi)$ denote the number of agents that rank x at position j in profile Π . Then, given a scoring rule \mathbf{s} , the score of alternative x is defined as

$$s_{\mathbf{CS}}(x, \Pi) = \sum_{j=1}^d \nu_j(x, \Pi) \cdot s_j.$$

We also assume that we have access to a set of *constraints* C that represents our (possibly partial) knowledge to an objective set of pairwise relations between the alternatives. Each constraint in C is given by an ordered pair of alternatives (x, y) , has a corresponding non-negative weight (of importance) $w(x, y)$, and requires that alternative x is ranked higher than alternative y in the outcome of the scoring rule \mathbf{s} . For a pair of alternatives (x, y) , let $\delta_j(x, y, \Pi) = \nu_j(x, \Pi) - \nu_j(y, \Pi)$. Now, observe that, in order for alternative x to be ranked above y with certainty in the final ranking, it must be $s_{\mathbf{CS}}(x, \Pi) > s_{\mathbf{CS}}(y, \Pi)$ and, equivalently,

$$\sum_{j=1}^d \delta_j(x, y, \Pi) \cdot s_j > 0.$$

Using $\boldsymbol{\delta}(x, y, \Pi) = (\delta_1(x, y, \Pi), \dots, \delta_d(x, y, \Pi))$, the above expression can be compactly written as the dot product $\boldsymbol{\delta}(x, y, \Pi) \cdot \mathbf{s} > 0$.

For our purposes, instead of thinking of a profile Π as the set of rankings provided by the agents, it is convenient to describe it using the quantities $\boldsymbol{\delta}(x, y, \Pi)$ for every constraint (x, y) in C ; we use the notation $\boldsymbol{\delta}(\Pi)$ to denote the set of these quantities and will simply refer to it as the profile.

Now, problem **OptPSR** (standing for “optimizing positional scoring rules”) is defined as follows. We are given a profile $\boldsymbol{\delta}(\Pi)$ and a set C of constraints. The goal of **OptPSR** is to find the scoring rule \mathbf{s} that produces a ranking of all alternatives so that the total weight (or gain)

$$\mathfrak{g}(\mathbf{s}, \boldsymbol{\delta}(\Pi), C) = \sum_{(x, y) \in C} w(x, y) \cdot \mathbb{1} \{ \boldsymbol{\delta}(x, y, \Pi) \cdot \mathbf{s} > 0 \},$$

of satisfied constraints is maximized. The quantity $\mathbb{1} \{X\}$ takes value 1 if X is true and 0 otherwise.

Let us now give an equivalent view of **OptPSR**. A scoring rule \mathbf{s} can be thought of as a point in \mathbb{R}^d , and, in particular, in the region R_0 of \mathbb{R}^d formed by the inequalities $s_i - s_{i+1} \geq 0$ for $i = 1, \dots, d-1$ and $s_d \geq 0$ that define all valid scoring vectors. We can define subregions of R_0 by considering any subset $C' \subseteq C$ of constraints and the inequality $\boldsymbol{\delta}(x, y, \Pi) \cdot \mathbf{s} > 0$ for every constraint associated with the pair of alternatives $(x, y) \in C'$ and the inequality $\boldsymbol{\delta}(x, y, \Pi) \cdot \mathbf{s} \leq 0$ for every constraint $(x, y) \in C \setminus C'$. In this way, the collection of all subsets of constraints in C partition R_0 into disjoint subregions (of course, some of them may be infeasible). Hence, in order to maximize $\mathfrak{g}(\mathbf{s}, \boldsymbol{\delta}(\Pi), C)$, it suffices to find any point \mathbf{s} in the non-empty subregion of R_0 that satisfies the subset of constraints with maximum total weight.

To do so, we can enumerate all subsets of constraints of C , check feasibility of the corresponding regions using linear programming, and report any point in the subregion that yields the highest gain.

This algorithm takes time polynomial in $2^{|C|}$ and d , assuming that it receives $\delta(\Pi)$ and C as input. In the next section, we will present an algorithm that uses a more clever enumeration of the feasible subregions in order to get the one that yields the maximum gain.

3 An improved OptPSR algorithm

We will present another (exact) OptPSR algorithm whose running time depends exponentially only on the parameter d and, hence, is polynomial when d is a constant.

The algorithm computes a pool of non-empty subregions of R_0 , each of which satisfies a different subset of constraints. Initially, the pool consists of region R_0 only and is updated as new constraints of C are considered. When a new constraint is considered, each region in the current pool can be split into two subregions consisting of the points that satisfy the constraint and those that do not satisfy it, respectively; the whole region is retained in the pool if all its points satisfy or (exclusively) do not satisfy the constraint.

In particular, the algorithm considers the constraints of C one by one. At each step t of the algorithm, a pool \mathcal{P} of regions is kept; at the beginning of each step, all regions in the pool are active. For each region R in \mathcal{P} , the algorithm keeps the gain $\text{val}(R)$ that is obtained by the constraints which have been considered until step t and are satisfied by scoring vectors of region R . The algorithms begin its execution having only region R_0 in the pool. When a new constraint (x, y) with weight $w(x, y)$ is considered, the algorithm attempts to update each active region R of \mathcal{P} as follows. It defines the candidate regions R^{xy} and $R^{\neg xy}$ such that

- R^{xy} is defined by the inequalities that form R together with inequality $\delta^{xy} \cdot \mathbf{s} > 0$ (that defines the set of points that satisfy constraint (x, y)), and
- $R^{\neg xy}$ is defined by the inequalities that form R together with inequality $\delta^{xy} \cdot \mathbf{s} \leq 0$ (that defines the set of points that do not satisfy constraint (x, y)).

If both R^{xy} and $R^{\neg xy}$ are non-empty (i.e., the corresponding sets of inequalities are feasible), the algorithm includes both R^{xy} and $R^{\neg xy}$ in \mathcal{P} as inactive, sets their gains $\text{val}(R^{xy}) := \text{val}(R) + w(x, y)$ and $\text{val}(R^{\neg xy}) := \text{val}(R)$, and removes region R from the pool. If only R^{xy} is feasible (and $R^{\neg xy}$ is infeasible), $\text{val}(R)$ is increased by $w(x, y)$. If only $R^{\neg xy}$ is feasible, the algorithm does nothing. In the last two cases, no new region is added to the pool. Clearly, it cannot be the case that both R^{xy} and $R^{\neg xy}$ are infeasible. Note that feasibility can be checked efficiently by solving linear programs with d variables and up to $|C|$ constraints. At the end of step t (i.e., when there is no other active region in the pool to be considered), the inactive regions become active and the algorithm proceeds with step $t + 1$.

When all constraints of C have been considered, the algorithm computes the active region R^* with maximum $\text{val}(R^*)$ and returns any scoring vector in R^* . An example of an execution of the algorithm with $d = 2$ is depicted in Figure 1.

Theorem 1. *Given an instance of OptPSR consisting of a set of constraints C and a profile $\delta(\Pi)$, the algorithm above correctly returns a solution in time $\mathcal{O}(|C|^d \cdot \text{poly}(|C|, d))$.*

Proof. The correctness of the algorithm should be apparent. It considers the whole space of points in \mathbb{R}^d which corresponds to scoring vectors and divides it into all (sub)regions defined for every inclusion-maximal subset of constraints that are satisfied simultaneously. Among all these regions, it finds the one with points that correspond to scoring vectors that satisfy constraints of C with maximum total weight.

Expanding R_0 into the regions in the pool when the last constraint of C is considered can be thought of as a non-complete binary tree T with nodes corresponding to regions (see Figure 1 for an example). T is rooted at a node corresponding to R_0 and is such that each node at level $t - 1$, corresponding to a region R , has two children at level t if the region R was split in and replaced by two subregions at

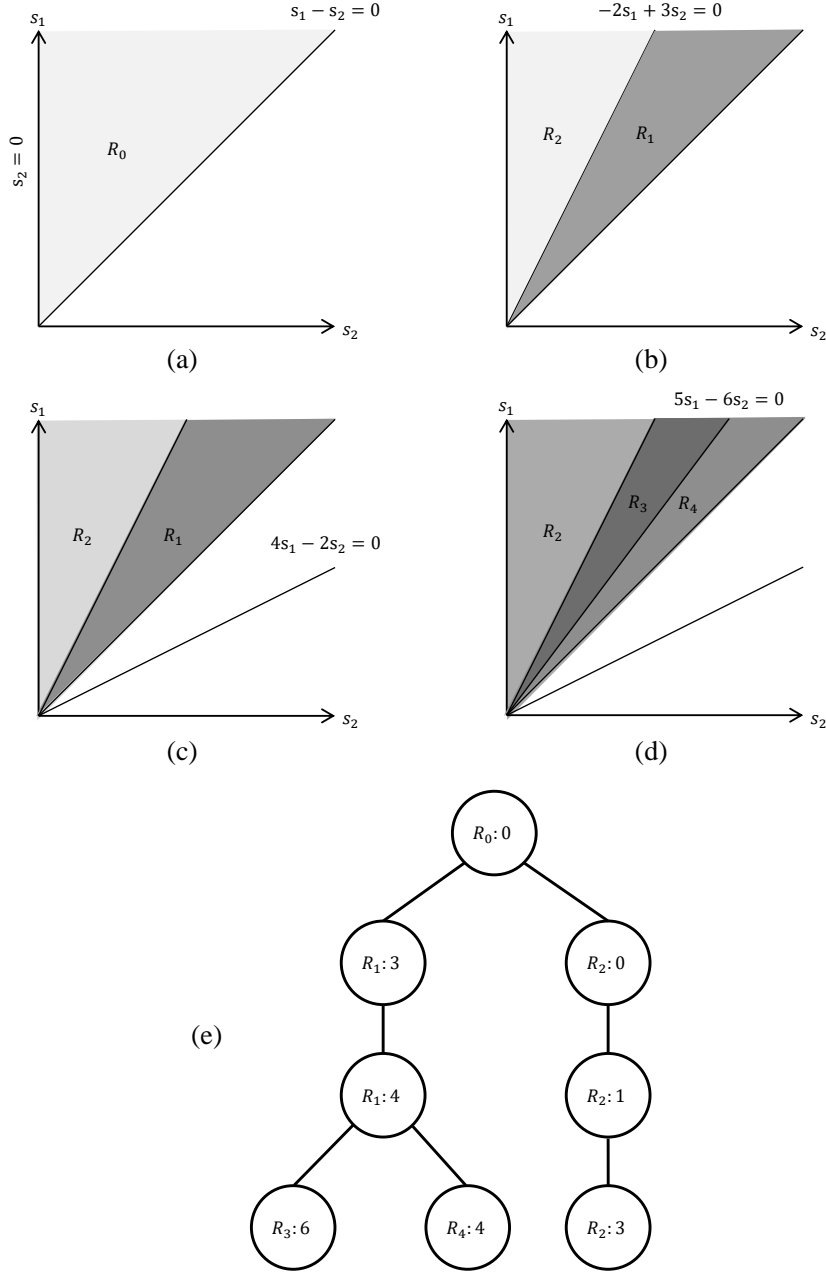


Figure 1: An example of how the algorithm works on a simple instance Π with $d = 2$. The set C has three constraints (x_1, y_1) , (x_2, y_2) , and (x_3, y_3) with corresponding weights 3, 1 and 2. The profile is such that $\delta(x_1, y_1, \Pi) = (-2, 3)$, $\delta(x_2, y_2, \Pi) = (4, -2)$, and $\delta(x_3, y_3, \Pi) = (5, -6)$. So, the constraints defined the inequalities $-2s_1 + 3s_2 > 0$, $4s_1 - 2s_2 > 0$, and $5s_1 - 6s_2 > 0$. (a) Initially, the algorithm has region R_0 in the pool. (b) At the next step, the algorithm considers constraint (x_1, y_1) and replaces R_0 with regions R_1 (which is the subregion of R_0 with $-2s_1 + 3s_2 > 0$ that satisfies the first constraint and has gain 3) and R_2 (with gain 0). (c) Next, constraint (x_2, y_2) leaves both regions R_2 and R_3 in the pool. (d) Finally, the third constraint (x_3, y_3) replaces region R_1 by regions R_3 and R_4 . In (e), the evolution of the content of the pool, together with the gains of the corresponding subregions are illustrated. The region with the maximum gain is R_3 and the algorithm will output some scoring vector from this region.

step t and has one child otherwise (indicating that the region was retained in the pool during step t). The total time required to find all regions is proportional to the size of T . Since all non-leaf nodes have at least one child, the size of T is at most its height $|C|$ times the number of leaves. The number of leaves is essentially the number of different non-empty regions, which is upper-bounded by the number of different sign patterns that the quantities $\delta^{xy} \cdot s$ define for each constraint (x, y) in C . Since these $|C|$ quantities are linear functions over the d coordinates of vector s , a result due to Alon [1] (see also Warren [17]) yields that the total number of different sign patterns is at most $\left(\frac{8e|C|}{d}\right)^d$. For each of the nodes of T , feasibility can be checked by solving two linear programs with d variables and at most $|C|$ constraints in time $\text{poly}(|C|, d)$. The theorem follows. \square

By Theorem 1, we obtain the following corollary. For comparison, the naive algorithm presented at the end of the previous section is polynomial in the very special case where $|C|$ is at most logarithmic in d .

Corollary 2. *The algorithm solves instances of OptPSR with constant d in polynomial time.*

4 Approximating OptPSR

As the running time of the exact algorithm of the previous section depends exponentially on d , our aim here is to design much faster (i.e., polynomial-time) algorithms that compute approximate OptPSR solutions. As we will see, an extremely simple scoring rule achieves a $1/d$ -approximation, i.e., the total weight of the constraints it satisfies is at least $1/d$ times the total weight satisfied by an optimal OptPSR solution. For $t \in [d]$, the t -approval rule is a positional scoring rule that uses the scoring vector that has 1 in the first t positions and 0 in the remaining ones.

Theorem 3. *For every instance of OptPSR with parameter d , there exists some $t \in [d]$ so that t -approval is a $1/d$ -approximate solution. This bound is tight.*

Proof. For the upper bound, consider a profile $\delta(\Pi)$ and set of constraints C . We partition the constraints of C into d disjoint sets C_1, C_2, \dots, C_d so that the t -th set is defined as

$$C_t = \left\{ (x, y) \in C : \sum_{k=1}^{\ell} \delta_k(x, y, \Pi) \leq 0 \text{ for all } \ell \leq t-1 \text{ and } \sum_{k=1}^t \delta_k(x, y, \Pi) > 0 \right\},$$

for $t = 1, 2, \dots, d$. Observe that the t -approval rule satisfies all constraints in the set C_t . Actually, set C_t is defined as the set of constraints that are satisfied by t -approval but not by ℓ -approval for $\ell < t$. Hence, the sets C_1, \dots, C_d are disjoint and there exists $t^* \in [d]$ (with \mathbf{t}^* being the scoring vector of t^* -approval) such that $\mathfrak{g}(\mathbf{t}^*, \delta(\Pi), C) \geq \sum_{(x,y) \in C_{t^*}} w(x, y) \geq \frac{1}{d} \sum_{(x,y) \in C} w(x, y)$. As the maximum possible gain cannot exceed $\sum_{(x,y) \in C} w(x, y)$, we have that t^* -approval is $1/d$ -approximate as desired.

For the lower bound, we will present an OptPSR instance such that any t -approval, with $t \in [d]$, is (at least) $1/d$ -approximate. The instance has d pairs of alternatives (x_t, y_t) as constraints with $w(x_t, y_t) = 1$ for $t \in [d]$. We will build a profile Π^* so that the t -approval rule satisfies only constraint (x_t, y_t) , while there exists a scoring rule that simultaneously satisfies all constraints. Consider quantities a_1, a_2, \dots, a_d with positive integer values such that $\sum_{j=1}^d \frac{1}{1+a_j} < 1$. The profile is defined as follows:

- Alternative x_1 appears a_1 times in position 1, and alternative y_1 appears $1 + a_1$ times in position 2. This means that $\delta_1(x_1, y_1, \Pi^*) = a_1$, $\delta_2(x_1, y_1, \Pi^*) = -1 - a_1$ and $\delta_j(x_1, y_1, \Pi^*) = 0$ for $j \geq 3$.

- For $2 \leq t \leq d-1$, alternative x_t appears $1 + a_t$ times in position t , and alternative y_t appears once in position 1 and $1 + a_t$ times in position $t+1$. This means that $\delta_1(x_t, y_t, \Pi^*) = -1$, $\delta_t(x_t, y_t, \Pi^*) = 1 + a_t$, $\delta_{t+1}(x_t, y_t, \Pi^*) = -1 - a_t$ and $\delta_j(x_t, y_t, \Pi^*) = 0$ for $j \notin \{1, t, t+1\}$.
- Alternative x_d appears $1 + a_d$ times in position d , and alternative y_d appears once in position 1. This means that $\delta_1(x_d, y_d, \Pi^*) = -1$, $\delta_d(x_d, y_d, \Pi^*) = 1 + a_d$ and $\delta_j(x_d, y_d, \Pi^*) = 0$ for $2 \leq j \leq d-1$.
- The rest of the positions in the votes are filled with additional alternatives that do not appear in the constraints.

Observe that, for $t \in [d]$, we have that $\sum_{j=1}^t \delta_j(x_t, y_t, \Pi^*) = a_t > 0$, $\sum_{j=1}^{t-1} \delta_j(x_t, y_t, \Pi^*) = -1$ and $\sum_{j=1}^\ell \delta_j(x_t, y_t, \Pi^*) = -1$ for $\ell > t$. Hence, the t -approval rule satisfies only constraint t for a total weight of 1.

Now we will show that there exists a scoring rule $\mathbf{s} = (s_1, s_2, \dots, s_d)$ that satisfies all constraints. Let $\epsilon > 0$ be some arbitrary small constant and consider the scoring vector \mathbf{s} with $s_1 = \frac{\epsilon \sum_{j=1}^d \frac{1}{1+a_j}}{1 - \sum_{j=1}^d \frac{1}{1+a_j}}$, $s_2 = \frac{a_1 s_1 - \epsilon}{1+a_1}$, and $s_i = s_{i-1} - \frac{s_1 + \epsilon}{1+a_{i-1}}$ for $i \geq 3$. First, observe that this is a valid scoring rule since, it is clear that $s_i \geq s_{i+1}$ for all $i \in [d-1]$ and, furthermore, it can be easily seen that $s_d \geq 0$ as well. Moreover, this scoring rule satisfies all constraints since $\sum_{j=1}^d \delta_j(x_t, y_t, \Pi^*) \cdot s_j = \epsilon$ for every $t \in [d]$. Hence, any t -approval is $1/d$ -approximate. \square

On the negative side, we show that our problem is not only hard, but also hard to approximate within some constant.

Theorem 4. *For every constant $\eta > 0$, OptPSR is hard to approximate within $23/24 + \eta$.*

Proof. We use a reduction from MAX-3LIN-2, the problem of maximizing the number of satisfied equations in an over-determined system of linear equations modulo 2. An instance of MAX-3LIN-2 consists of n binary variables $x_i \in \{0, 1\}$ and m equations of the forms $x_i \oplus x_j \oplus x_k = 0$ and $x_i \oplus x_j \oplus x_k = 1$, where \oplus denotes addition modulo 2 and its objective is to find an assignment to the variables so that the number of satisfied equations is maximized. Below, we use the term α -equation to refer to an equation of the form $x_i \oplus x_j \oplus x_k = \alpha$ (for $\alpha \in \{0, 1\}$).

Given an instance of MAX-3LIN-2, our reduction constructs in polynomial-time an instance of OptPSR that has a scoring rule that satisfies constraints of total weight $11m + L$ if and only if the MAX-3LIN-2 instance has an assignment satisfying L equations. A famous result by Håstad [9] states that it is hard to distinguish in time polynomial in n and m whether a given instance of MAX-3LIN-2 has an assignment that satisfies at least $(1 - \eta')m$ equations or any assignment satisfies at most $(1/2 + \eta')m$ equations, for any constant $\eta' > 0$. As a consequence of our reduction, we obtain that it is hard to distinguish between instances of OptPSR that have a scoring rule that satisfies constraints of total weight at least $(12 - \eta')m$ and instances of OptPSR in which the total weight of the constraints satisfied by any scoring rule is at most $(23/2 + \eta')m$. An inapproximability bound of $23/24 + \eta$ (for every constant $\eta > 0$) then follows by standard arguments.

Without loss of generality, we can assume that the scoring vectors $\mathbf{s} = (s_1, s_2, \dots, s_d)$, that we seek for, have $s_1 = d$ and the remaining scores are defined in terms of $d-1$ variables $a_1, a_2, \dots, a_{d-1} \geq 0$ as $s_{i+1} = s_i - a_i$ (or, consequently, $s_j = d - \sum_{k=1}^{j-1} a_k$) for $i = 1, \dots, d-1$ so that $\sum_{i=1}^{d-1} a_i \leq d$. Hence, a constraint (y, z) requiring that the score of y is higher than the score of z can be expressed as a linear inequality of the variables a_j with $j \in [d-1]$. The assumption that $s_1 = d$ allows for inequalities that have non-zero constant terms. We define linear inequalities corresponding to constraints first; later, we also construct the profile and specify the constraints as pairs of alternatives and corresponding weights that are consistent to these linear inequalities. Let m_i be the number of equations in which variable x_i

participates. We set ϵ to be a small constant such that $0 < \epsilon \leq 1/d$ and $1/\epsilon$ is an integer. The instance of OptPSR can be expressed with the following inequalities that represent constraints:

- For every variable x_i , we have the four inequalities $a_i > 0$, $a_i < \epsilon$, $a_i > 1$ and $a_i < 1 + \epsilon$ of weight m_i each.
- For every equation, there are four inequalities of unit weight each:
 - if the equation is of the form $x_i \oplus x_j \oplus x_k = 0$, the inequalities are $a_i + a_j + a_k > 0$, $a_i + a_j + a_k < \epsilon$, $a_i + a_j + a_k > 2$ and $a_i + a_j + a_k < 2 + \epsilon$, and
 - if the equation is of the form $x_i \oplus x_j \oplus x_k = 1$, the inequalities are $a_i + a_j + a_k > 1$, $a_i + a_j + a_k < 1 + \epsilon$, $a_i + a_j + a_k > 3$ and $a_i + a_j + a_k < 3 + \epsilon$.

These inequalities are implemented as follows. Let $d = n + 1$. For every variable x_i , with $i \in [n]$, we have four constraints (y_i^t, z_i^t) , where $t \in \{1, 2, 3, 4\}$, of weight m_i each. In the profile, alternatives y_i^t and z_i^t appear in specific positions as follows:

- Alternative y_i^1 appears once in position i , and alternative z_i^1 appears once in position $i + 1$. Then, the constraint (y_i^1, z_i^1) corresponds to the inequality $s_i - s_{i+1} > 0$ or, equivalently, $a_i > 0$.
- Alternative y_i^2 appears once in position 1 and d/ϵ times in position $i + 1$, and alternative z_i^2 appears d/ϵ times in position i . The constraint (y_i^2, z_i^2) corresponds to the inequality $s_i - s_{i+1} < \epsilon$ or $a_i < \epsilon$.
- Alternative y_i^3 appears d times in position i , and alternative z_i^3 appears once in position 1 and d times in position $i + 1$. The constraint (y_i^3, z_i^3) corresponds to the inequality $s_i - s_{i+1} > 1$ or $a_i > 1$.
- Alternative y_i^4 appears $1/\epsilon + 1$ times in position 1 and d/ϵ times in position $i + 1$, and alternative z_i^4 appears d/ϵ times in position i . The constraint (y_i^4, z_i^4) corresponds to the inequality $s_i - s_{i+1} < 1 + \epsilon$ or $a_i < 1 + \epsilon$.

Observe that three of the four inequalities corresponding to variable x_i can simultaneously be satisfied when $a_i \in (0, \epsilon) \cup (1, 1 + \epsilon)$, and only two of them are satisfied for any other value of a_i .

For every equation $\ell \in [m]$, we have four constraints (b_ℓ^t, c_ℓ^t) , where $t \in \{1, 2, 3, 4\}$, of unit weight each. In the profile, these alternatives appear in specific positions depending on whether equation ℓ is a 0- or a 1-equation. In the case where it is a 0-equation of the form $x_i \oplus x_j \oplus x_k = 0$, we have:

- Alternative b_ℓ^1 appears once in positions i, j and k , and alternative c_ℓ^1 appears once in positions $i + 1, j + 1$ and $k + 1$. Then, the constraint (b_ℓ^1, c_ℓ^1) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} > 0$ or, equivalently, $a_i + a_j + a_k > 0$.
- Alternative b_ℓ^2 appears d times in position 1 and d/ϵ times in positions $i + 1, j + 1$ and $k + 1$, and alternative c_ℓ^2 appears d/ϵ times in positions i, j and k . The constraint (b_ℓ^2, c_ℓ^2) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} < \epsilon$ or $a_i + a_j + a_k < \epsilon$.
- Alternative b_ℓ^3 appears d times in positions i, j and k , and alternative c_ℓ^3 appears 2 times in position 1 and d times in positions $i + 1, j + 1$ and $k + 1$. The constraint (b_ℓ^3, c_ℓ^3) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} > 2$ or $a_i + a_j + a_k > 2$.
- Alternative b_ℓ^4 appears $2/\epsilon + 1$ times in position 1 and d/ϵ times in positions $i + 1, j + 1$ and $k + 1$, and alternative c_ℓ^4 appears d/ϵ times in positions i, j and k . The constraint (b_ℓ^4, c_ℓ^4) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} < 2 + \epsilon$ or $a_i + a_j + a_k < 2 + \epsilon$.

Observe that three of the inequalities corresponding to a 0-equation can simultaneously be satisfied when $a_i + a_j + a_k \in (0, \epsilon) \cup (2, 2 + \epsilon)$; otherwise, exactly two inequalities are satisfied. In the case where equation ℓ is a 1-equation of the form $x_i \oplus x_j \oplus x_k = 1$, we have:

- Alternative b_ℓ^1 appears d times in positions i, j and k , and alternative c_ℓ^1 appears once in position 1 and d times $i + 1, j + 1$ and $k + 1$. Then, the constraint (b_ℓ^1, c_ℓ^1) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} > 1$ or, equivalently, $a_i + a_j + a_k > 1$.
- Alternative b_ℓ^2 appears $1/\epsilon + 1$ times in position 1 and d/ϵ times in positions $i + 1, j + 1$ and $k + 1$, and alternative c_ℓ^2 appears d/ϵ times in positions i, j and k . The constraint (b_ℓ^2, c_ℓ^2) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} < 1 + \epsilon$ or $a_i + a_j + a_k < 1 + \epsilon$.
- Alternative b_ℓ^3 appears d times in positions i, j and k , and alternative c_ℓ^3 appears 3 times in position 1 and d times in positions $i + 1, j + 1$ and $k + 1$. The constraint (b_ℓ^3, c_ℓ^3) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} > 3$ or $a_i + a_j + a_k > 3$.
- Alternative b_ℓ^4 appears $3/\epsilon + 1$ times in position 1 and d/ϵ times in positions $i + 1, j + 1$ and $k + 1$, and alternative c_ℓ^4 appears d/ϵ times in positions i, j and k . The constraint (b_ℓ^4, c_ℓ^4) corresponds to the inequality $s_i - s_{i+1} + s_j - s_{j+1} + s_k - s_{k+1} < 3 + \epsilon$ or $a_i + a_j + a_k < 3 + \epsilon$.

Again, for every 1-equation, we have three inequalities that can simultaneously be satisfied when $a_i + a_j + a_k \in (1, 1 + \epsilon) \cup (3, 3 + \epsilon)$; otherwise, exactly two inequalities are satisfied.

In order for this profile to be valid, we use sufficiently many agents and additional alternatives as placeholders, so that the alternatives mentioned above have the appropriate number of appearances in the rankings.

We now prove that there exists a variable assignment for the MAX-3LIN-2 instance that satisfies L of its equations if and only if there exists a scoring rule that satisfies constraints of total weight $11m + L$. As we have discussed above, this is enough to complete the proof.

Consider an assignment that satisfies L of the equations. Then, the scoring rule defined by setting $a_i = x_i + \epsilon/4$ for $i \in [d - 1]$ satisfies:

- three out of the four inequalities corresponding to any variable x_i , since $a_i = \epsilon/4 \in (0, \epsilon)$ when $x_i = 0$ and $a_i = 1 + \epsilon/4 \in (1, 1 + \epsilon)$ when $x_i = 1$;
- three out of the four inequalities corresponding to any satisfied 0-equation $x_i \oplus x_j \oplus x_k = 0$ since $a_i + a_j + a_k = 3\epsilon/4 \in (0, \epsilon)$ when $x_i + x_j + x_k = 0$ and $a_i + a_j + a_k = 2 + 3\epsilon/4 \in (2, 2 + \epsilon)$ when $x_i + x_j + x_k = 2$;
- two out of the four inequalities corresponding to any unsatisfied 0-equation since $a_i + a_j + a_k \notin (0, \epsilon) \cup (2, 2 + \epsilon)$ in that case;
- three out of the four inequalities corresponding to any 1-equation $x_i \oplus x_j \oplus x_k = 1$ since $a_i + a_j + a_k = 1 + 3\epsilon/4 \in (1, 1 + \epsilon)$ when $x_i + x_j + x_k = 1$ and $a_i + a_j + a_k = 3 + 3\epsilon/4 \in (3, 3 + \epsilon)$ when $x_i + x_j + x_k = 3$;
- two out of the four inequalities corresponding to any unsatisfied 1-equation since $a_i + a_j + a_k \notin (1, 1 + \epsilon) \cup (3, 3 + \epsilon)$ then.

Hence, the total weight of the constraints satisfied is $3 \sum_{i=1}^n m_i + 3L + 2(m - L) = 11m + L$ since $\sum_{i=1}^n m_i = 3m$ due to the fact that all equations have 3 variables and the sum accounts for the total number of appearances of all variables.

Conversely, assume that we are given a scoring rule that satisfies constraints of total weight $11m + L$; we will show that there exists an assignment to the variable of the MAX-3LIN-2 instance that satisfies

L equations. First, we show that we can transform the scoring rule into a (possibly) different one with $a_i = \epsilon/4$ or $a_i = 1 + \epsilon/4$ for $i \in [d-1]$, without decreasing the total weight of the satisfied constraints.

For a variable $a_i \notin (0, \epsilon) \cup (1, 1 + \epsilon)$ we have that the satisfied inequalities are the following: exactly two out of the four variable inequalities and at most three out of the four inequalities for each of the m_i equations in which the variable x_i appears. This gives a weight of at most $2m_i + 3m_i = 5m_i$. By setting $a_i := \epsilon/4$, exactly three out of the four variable inequalities and at least two out of the four equation inequalities in which x_i appears are satisfied, for a total weight of at least $5m_i$. Clearly, there is no loss in weight after this change in the value of a_i .

Now, we slightly modify the variable values as follows: for all variables $a_i \in (0, \epsilon)$ we set $a_i := \epsilon/4$ and for all variables $a_i \in (1, 1 + \epsilon)$ we set $a_i := 1 + \epsilon/4$. The set of inequalities containing a_i that were satisfied before the modification are still satisfied after the update as well. This is trivial for the variable inequalities. For an equation inequality of the form $a_i + a_j + a_k < \alpha + \epsilon$ (respectively, $a_i + a_j + a_k > \alpha$) that was satisfied before the modification, at most α (respectively, at least α) of the three variables have values in $(1, 1 + \epsilon)$ before the modification. Clearly, the inequality is satisfied after the modification as well.

So, we can assume that we have total weight of $11m + L$ from satisfied constraints with the variables a_i taking values in $\{\epsilon/4, 1 + \epsilon/4\}$. Hence, $3 \sum_{i \in [n]} m_i = 9m$ comes as weight from satisfied variable inequalities (with three satisfied inequalities per variable). Then, the remaining weight comes from $2m + L$ satisfied equation inequalities. The definition of the reduction implies that there exist L equations in the MAX-3LIN-2 instance so that three among the four corresponding inequalities are satisfied. Then, it is easy to inspect that, if three among the four equation inequalities are satisfied when variables take values in $\{\epsilon/4, 1 + \epsilon/4\}$, then the assignment $x_i := a_i - \epsilon/4$ satisfies their corresponding equation as well. This yields an assignment with (at least) L satisfied equations and the proof is complete. \square

5 Experiments

We have conducted experiments for two different scenarios; we refer to them as *ppl* and *col*. In these two scenarios, we used as alternatives 48 countries and 36 cities, respectively. In both cases, the alternatives were used to define 392 different sets consisting of six alternatives each. The alternatives have been distributed to the different sets almost uniformly; each country appears in at least 47 and at most 52 sets and each city appears in at least 57 and at most 70 sets.

We used both real-world and synthetic data. Real-world data were collected as input from 392 participants in a technology exhibition at our home institution. Each of them was given two distinct sets of six countries and six cities. They were asked to rank the countries in terms of their population and the cities in terms of their cost of living. Synthetic data were obtained by simulating 392 agents who use the Plackett-Luce and Bradley-Terry noise models in order to compute random rankings.

The Bradley-Terry model [3] (BT, in short) is used by an agent in order to decide relations between all pairs of alternatives in her set as follows. Consider a pair of alternatives (x, y) with corresponding utilities (populations or cost of living indices) u_x and u_y . The agent decides to rank x above y with probability $\frac{u_x}{u_x + u_y}$ and y above x with probability $\frac{u_y}{u_x + u_y}$. If the relative ranks of all pairs of alternatives (that have been computed separately) do not define a ranking, the whole process is repeated.

In the Plackett-Luce model [11, 12] (PL, in short), an agent decides the ranking of the alternatives in her set sequentially. Let B be the set of alternatives the agent has to rank. Starting from the first position, the next undetermined position in the ranking is filled by alternative $x \in B$ with probability $\frac{u_x}{\sum_{y \in B} u_y}$. After a random selection, the chosen alternative is removed from B and the process continues for the next undetermined position and the remaining alternatives until all positions are filled.

The set of constraints were defined using population data for the 48 countries from `en.wikipedia.org` and cost of living index data for the 36 cities from `numbeo.com` as retrieved

in April 2016. In particular, in the ppl scenario, we have a constraint for each pair of countries x and y so that x is more populous than y . We consider two different weightings for constraints using weight that is either 1 or equal to the population difference between countries x and y . Unit weights are used when we care only about maximizing the number of correctly recovered population comparison between countries. However, there might be pairs that are really important to recover correctly, while some others are not that important. For example, it is important to conclude that China is ranked above Switzerland (their population difference is 1.3 billions) but an error in the comparison between Cuba and Belgium (both with populations around 11 millions) would not be that severe. Analogously, in the col scenario, we have a constraint for every pair of cities x and y so that x has higher cost of living index than y . The weight of the corresponding constraint is either 1 or equal to the cost of living index difference between the two cities.

Since all the profiles we experimented with have $d = 6$, one would expect that the exact algorithm presented in Section 3 would be the obvious choice in order to come up with the optimal scoring rule. Unfortunately, for the size of OptPSR instances we considered (with 1128 constraints for ppl and 630 constraints for col), this algorithm turned out to be really slow, even after implementing several heuristics that yield minor improvements to performance. This rather disappointing outcome together with the fact that d is small, forced us to consider scoring vectors with discretized scores (e.g., which are multiples of 0.05 or 0.02) in order to come up with approximations of the optimal scoring rule. This approach has yielded the vectors $(1, 0.5, 0.35, 0.2, 0.15, 0.05)$ and $(1, 0.65, 0.65, 0.35, 0.3, 0.25)$ for the ppl profile with unweighted and weighted constraints and the vectors $(1, 0.9, 0.3, 0.3, 0.24, 0)$ and $(1, 0.68, 0.68, 0.5, 0.22, 0.22)$ for the col profile with unweighted and weighted constraints, respectively.

We compare the optimal OptPSR solution (obtained in this way) to several well-known scoring rules such as the Borda count (defined by the scoring vector $(d, d - 1, \dots, 1)$), the harmonic rule (also known as Dowdall; defined by the scoring vector $(1, 1/2, \dots, 1/d)$), and t -approval rules. Tables 1 and Table 2 show the performance of these scoring rules in all the OptPSR instances we experimented with.

In Table 1, which contains data for instances with unweighted constraints, we observe that Borda and Harmonic outperform all approval rules in all cases besides the col scenario with PL agents, where 4-approval is slightly better than Harmonic. Also, there are cases (e.g., ppl profile with BT agents) where Borda performs better than Harmonic and vice versa (e.g., see the results for real-world data). In all scenarios, the values of Borda, Harmonic and the best approval rule have an average difference of 2.3%, 2.36% and 3.1%, respectively, from the optimal values (with maximum difference values of 3.43%, 5.07% and 4.32% that are all observed for the col profile with PL agents). Interestingly, approval rules have amazingly better performance than what their worst-case analysis from Theorem 3 indicates.

Clearly, Table 2 shows significantly better results from (almost) all scoring rules on OptPSR instances with weighted constraints. This is to be expected since solutions improve significantly when heavy pairwise relations are correctly recovered. Here, Borda, Harmonic and the best approval rule are all closer to the optimal performance. Now, the average distance is 1.05%, 1.08%, and 1.2%; again, the largest differences are observed for the col profile with PL agents.

6 Open problems

Our work reveals several open problems. First, we would like to determine the approximability of OptPSR. Is there a polynomial time algorithm with constant approximation ratio? Second, we would like to design an exact algorithm that is practical. Our ambitious goal is to be able to solve OptPSR instances like the ones we used in our experiments. Third, we would like to analyze theoretically scoring rules in random profiles that have been produced by Plackett-Luce or Bradley-Terry agents. Also, considering agents following other noise models that are close to real-world agents definitely deserves investigation.

rule	real data		synthetic (BT)		synthetic (PL)	
	ppl	col	ppl	col	ppl	col
opt.	81.83	83.97	94.54	93.74	93.19	88.20
borda	79.87	81.43	93.16	91.03	91.59	84.67
harm.	80.94	82.54	92.84	91.34	90.50	83.13
1-app.	77.75	78.09	83.69	87.89	83.90	75.72
2-app.	78.19	79.36	89.71	90.65	88.72	81.29
3-app.	79.43	80.48	91.69	90.93	90.30	83.73
4-app.	77.57	79.68	90.00	89.44	89.70	83.88
5-app.	73.14	72.86	81.08	84.45	82.83	80.66
6-app.	32.53	51.43	32.54	51.43	32.54	51.43

Table 1: Performance (as percentage of the total weight of all constraints) of scoring rules on instances with unweighted constraints. For synthetic profiles with BT and PL agents, the simulation was repeated 500 times; the values shown here are averages.

rule	real data		synthetic (BT)		synthetic (PL)	
	ppl	col	ppl	col	ppl	col
opt.	95.98	92.93	99.66	98.69	99.48	96.02
borda	94.56	91.57	99.49	97.66	99.23	93.96
harm.	95.42	92.01	99.42	97.80	99.02	92.58
1-app.	94.85	89.84	98.59	96.48	97.86	86.34
2-app.	95.24	90.40	99.29	97.67	98.86	91.50
3-app.	93.68	90.83	99.05	97.70	99.04	93.37
4-app.	92.63	90.06	97.46	96.91	98.38	93.51
5-app.	84.61	82.00	87.94	93.81	91.69	90.99
6-app.	39.66	57.04	39.88	56.96	39.88	56.96

Table 2: Performance of scoring rules on instances with weighted constraints. Again, for synthetic profiles with BT and PL agents, the values indicate average performance from 500 simulations.

References

- [1] N. Alon. Tools from higher algebra. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 2, pages 1749–1783. MIT Press, 1996.
- [2] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [3] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [4] F. Brandt, V. Conitzer, U. Endriss, J. Lang, H. Moulin, and A. D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [5] I. Caragiannis, G. A. Krimpas, and A. A. Voudouris. Aggregating partial rankings with applications to peer grading in massive online open courses. In *Proceedings of the 14th International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*, pages 675–683, 2015.

- [6] I. Caragiannis, G. A. Krimpas, and A. A. Voudouris. How effective can simple ordinal peer grading be? In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 323–340, 2016.
- [7] M. M. de Weerd, E. H. Gerding, and S. Stein. Minimising the rank aggregation error. In *Proceedings of the 15th International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*, pages 1375–1376, 2016.
- [8] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference (WWW)*, pages 613–622, 2001.
- [9] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [10] E. Law and L. von Ahn. *Human computation*. Morgan & Claypool Publishers, 2011.
- [11] R. D. Luce. *Individual choice behavior: A theoretical analysis*. Wiley, 1959.
- [12] R. L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 24(2):193–202, 1975.
- [13] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. The learnability of voting rules. *Artificial Intelligence*, 173(12-13):1133–1149, 2009.
- [14] K. Raman and T. Joachims. Methods for ordinal peer grading. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1037–1046, 2014.
- [15] D. Sculley. Rank aggregation for similar items. In *Proceedings of the 7th SIAM International Conference on Data Mining (SDM)*, pages 587–592, 2007.
- [16] N. B. Shah, J. K. Bradley, A. Parekh, M. Wainwright, and K. Ramchandran. A case for ordinal peer-evaluation in MOOCs. In *Neural Information Processing Systems (NIPS): Workshop on Data Driven Education*, 2013.
- [17] H. Warren. Lower bounds for approximation by non-linear manifolds. *Transaction of the American Mathematical Society*, 133:167–178, 1968.